

Building a Knowledge Graph for Recommending Experts

Behnam Rahdari
ber58@pitt.edu
University of Pittsburgh
Pittsburgh, PA

Peter Brusilovsky
peterb@pitt.edu
University of Pittsburgh
Pittsburgh, PA

ABSTRACT

Recommending experts is an important challenge in many contexts. In this paper, we present a method to build a knowledge graph by integrating data from Google Scholar and Wikipedia to help students find a research advisor or thesis committee member. This knowledge graph is used to power the exploratory search interface to recommend similar keywords and relevant scholars to the students with a limited level of knowledge and familiarity with the subject of research.

CCS CONCEPTS

• **Information systems** → **Graph-based database models; Information integration; Web searching and information discovery.**

KEYWORDS

Data Integration, Knowledge Graph, Recommender Systems

ACM Reference Format:

Behnam Rahdari and Peter Brusilovsky. 2019. Building a Knowledge Graph for Recommending Experts. In *KI2KG '19: 1ST INTERNATIONAL WORKSHOP ON CHALLENGES AND EXPERIENCES FROM DATA INTEGRATION TO KNOWLEDGE GRAPHS*, August 05, 2019, Anchorage, AK. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Recommending experts is an important challenge in many contexts. The nature of this challenge is to find a knowledgeable person with an advance expertise in one or more target topics among a large number of potential candidates. A well-explored examples are finding an expert for a specific project within a large company or finding a doctor with advance knowledge of a specific disease in a large city. While in these two contexts large companies and hospitals use knowledge management techniques to catalogue key areas of expertise and use it to represent information about each expert, finding experts in other contexts could be more challenging.

The context that we target out in this paper is finding a research advisor. Every year many undergraduate, master-level and PhD students are facing a difficult of finding a research advisor. While large universities have many highly knowledgeable faculty, finding

one with the expertise that matches student interests, requirements, and preparation is a hard problem. Whether this search is for finding an advisor for a summer research project, a faculty sponsor for an independent study, or a committee member for a PhD thesis, online sources frequently fail the students and they resort to the word of mouth within a limited circle of instructors, classmates, and university staff. One problem is a wide variety of sources where advisor information could be found online (department directories, publication sites, funding agency pages, personal home pages, etc.) Each of these sources covers only some aspects of faculty expertise and frequently represent only a subset of available advisors. Another problem is the lack of "expertise catalog". A university usually offers a catalog of courses and majors, but not a fine-grained catalog of expertise areas covered by faculty. As a result students frequently can't even properly name their target area of interest or formulate a Web search query when looking for advisors.

The focus of our project is to offer a single-access-point exploratory search system, which allows students to discover their target areas of interest and find relevant advisors within these areas. In its core, the platform uses a knowledge expertise graph, which represents multiple connections between research topics and prospective research advisors within a large university or a large research field. We build this graph by processing several knowledge sources about faculty and their research interests. This paper briefly reviews the kind of knowledge graph we built, the process of its building by information extraction, and the information exploration system powered by this knowledge.

2 BACKGROUND

The attempts to built "a map of science" representing most important areas of research expertise and their connections with experts have been done in the past, however, the lack of proper information sources made it hard to produce maps that are suitable for finding advisors. A good examples are attempts presented in [9] and [6], where academic journals are used as proxies of expertise areas and a map of science is built by clustering journals by co-publication links. While this map is useful as a "big picture" of science, its use for advisor finding is problematic since it represents expertise on a very coarse-grain level and misses many prospective advisor who are not frequent journal authors. However, the emergence of modern sites powered by a combination of advanced information processing and collective wisdom makes the task of building a fine-grain knowledge network of experts and expertise areas feasible. In our work we rely most extensively to two of these sites - *Google Scholar* and *Wikipedia*.

Google Scholar has been long recognized among the best freely accessible academic information sources in term of coverage and accessibility [3]. It has been compared positively with number of similar citation services namely *Web of Science* [1], *PubMed*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KI2KG '19, August 05, 2019, Anchorage, AK

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

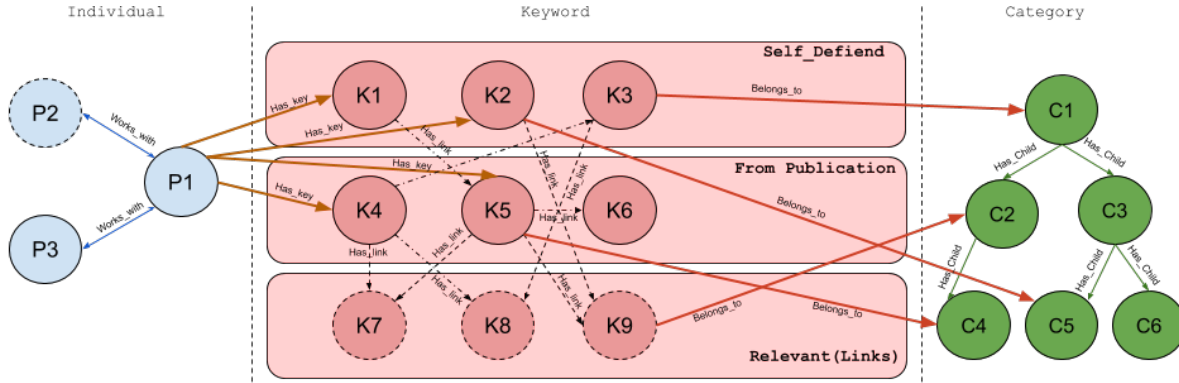


Figure 1: High level Knowledge Graph schema: from left to right, "individual nodes" (blue) store scholar’s demographic information , "keyword nodes" (red) keep detailed information about topic keywords and "category nodes"(green) convey the hierarchical association between categories and semantic relationship between keywords.

and *Scopus* [2, 5]. Yet, although *Google Scholar* contains nearly 160 million documents [7] covering a large portion of published documents, the lack of semantic connections between concepts and keywords within these documents makes it difficult to use the system for finding advisors, especially by less experienced students.

Wikipedia is commonly used by researchers to compute the semantic relatedness within documents [8, 10, 11], extract Open Information [13] and mine meaning using relations, facts and description to extracts and makes use of the concepts [4].

3 BUILDING THE KNOWLEDGE GRAPH

To support students in finding advisors we created a knowledge graph using data from *Google Scholar* and enriching it semantically using *Wikipedia*. In turn, this graph was used to power an interactive exploratory recommendation interface, which make the task of advisor-finding easier, especially for students with a limited level of knowledge and familiarity with the subject of research. To support several advisor-finding scenarios, we build several versions of the graph. The graph presented in this paper is focused on task of finding a top expert in a specific topic of interest within some broad field of research (such as Artificial Intelligence) across many universities. This is a typical task for a student selecting a PhD program to join or for a senior PhD student looking for an external thesis committee member.

3.1 Data Sources

3.1.1 Google Scholar. We exploited the information of 1000 active scholars in two popular fields of computer science; *Artificial Intelligent* and *Computer Architecture* (focusing on top 500 scholars per field). For each individual, we extracted the following information (see Table 1):

- **Name** : Full name of the scholar.
- **Affiliation** : The university or research institution the scholar affiliated with.
- **Verified Email Domain**: It used to check the validity of the scholar profile.
- **Self-Defined Keywords**: A list of up to five keywords defined by scholars to describe their research interests.

- **Citations** : The total number of citations received by all the scholar’s publications.
- **h-index** : The h-index measures the citation impact and productivity of a scholar’s publications. We use this measure alongside with other quantitative scores to re-order the results of the recommendations.
- **i10-Index** : i10-Index describes as the total number of scholar’s publications with 10 citations and more. This score, which is only used by *Google Scholar*, also used to re-rank the results of the recommendations.
- **Recent publications (20)** : We used 20 most recent publications to generate additional keywords representing current interests of each scholar. The keywords were extracted from the titles of recent publications as follows: After removing stop-words, we generated all the possible keyword candidates as uni-grams, bi-grams, and tri-grams. Next, we only kept the keywords that have an entry in *Wikipedia* (see keyword verification below).
- **Top Co-Authors (10)** : For each scholar, we extracted a list of top 10 co-authors from their *Google Scholar* profile.

3.1.2 Wikipedia. We used *Wikipedia* to add a semantic layer to profiles extracted from *Google Scholar*. Throughout this process, we also obtained some useful information that led to a stronger connection between keywords and enables us to add weight to each scholar-keyword relation. *Wikipedia* API has been used for the following purposes:

- **Keyword verification** : As mentioned before, we collected two sets of keywords for each scholar: self-defined keywords and keywords extracted from recent publications. *Wikipedia* API has been used to verify the validity of these keywords by using fuzzy match technique to find the *Wikipedia* entry describing the keyword. We removed all keywords that did not match with any article in *Wikipedia*. While *Wikipedia* might miss articles for some less popular research topics, we need to have all topic keywords explainable for the student audience and a match to a *Wikipedia* article was the

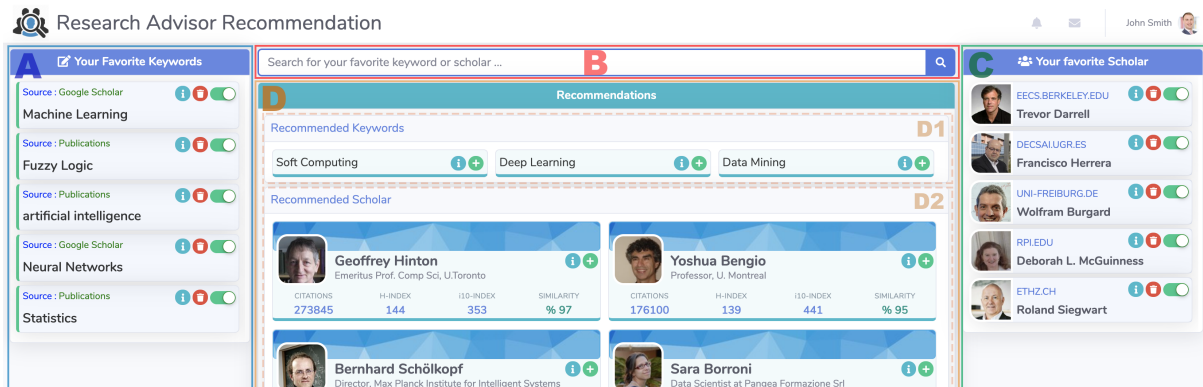


Figure 2: Interface Design of Research Advisor Recommendation

best way to assure it. For all remaining keywords, we calculated association weight between a keyword and a scholar as cosine similarity between the full-text Wikipedia entry of each keyword and concatenated text from scholar's recent publications.

- **Entry Summary** : To offer student users a short description of each topic keyword, we collected page summary for all keywords using Wikipedia API.
- **Top relevant keywords (10)** : Most (if not all) Wikipedia pages have multiple links to similar or related articles. We collected the top 10 links based on the number of their occurrences in each page. We employ these links to create a highly connected network of keywords.
- **Entry Categories** : Wikipedia uses categories to group similar articles. We extracted all categories associated with a page and used a full Wikipedia category hierarchy schema to find relationships between categories in our data set.

3.2 Graph Representation

We used Neo4j [12] graph database to represent information about all scholars and keywords. The overall schema of the knowledge graph is represented in Figure 1. As is shown, there are three distinct node types in the graph:

- **Individual** : This node type conveys demographic information about scholars including full name, affiliation, verified email domain, and URLs of personal homepage and Wikipedia page (if exist). *Individual* nodes are connected via "works_with" links which represent the co-authorship relations between scholars. An *individual* node also connects to several *keyword* nodes that represent the scholar's research interests and expertise. The *individual* nodes with a dashed border represent scholars added via *co-authorship extraction* who are not among the top 500 extracted scholars. These nodes are not considered in the final recommendations and only used to indicate the connections of the top scholars.
- **Keyword** : There are three types of *keyword* nodes. Self-Defined keywords, keywords extracted from recent scholar's publications, and relevant keywords that are represent the connection between two other types (shown by a dashed border) and will not appear in the recommendations. The

relationship between keywords represented by "has_link" arc is established if the target node has been mentioned in the source node's entry page. *Keyword* nodes are connected to *individual* nodes via "has_key" and to *category* nodes via "belongs_to" arcs.

- **Category** : We employed a full hierarchical schema of Wikipedia categories to represent the inter-connectivity between categories in our data set. These relations are presented as "has_child" arc in Figure 1. The *category* nodes are used to find the semantic relationship between keywords.

4 USING THE KNOWLEDGE GRAPH

4.1 Interface Design

The knowledge graph is used to power the exploratory search interface for finding advisors. the interface consist of four main sections.

4.1.1 Instant search box. User can use the search box to search for topic keyword or scholars of interest (Figure 2:B). When the user starts typing, a list of matching keywords and scholars will appear. When an item is selected from the list, it will automatically added to proper location on left or right side of the interface. at the same time, an updated list of recommendations will be presented to the user.

4.1.2 Favorite keywords. This section (Figure 2:A) shows user's "favorite" keywords. User add keywords to this list using the instant search box or by clicking on the *plus button* next to each recommended keyword. User can interact with three buttons on the right side of each keywords to (1) see more information about that keyword (including Wikipedia summary, similar keywords and other scholars with this research interest), (2) remove the keyword from the favorite list and (3) enable/disable the effect of this keyword in the list of recommendations.

4.1.3 Favorite scholars. Similar to *favorite keywords*, the user can assemble a list of favorite scholars (Figure 2:C). A new favorite scholar could be added to the list from the instant search results or a list or recommended scholars by clicking on *plus button* next to a recommended scholar. The three buttons on the right side of each favorite scholars could be used to obtain more details about

| | Artificial Intelligent | | Computer Architecture | |
|------------------------|------------------------|--------|-----------------------|--------|
| | all | unique | all | unique |
| Self-defined Keywords | 1916 | 628 | 1671 | 493 |
| Publication Concepts | 5946 | 1985 | 5889 | 1650 |
| Relevant Keywords | 37339 | 24712 | 30677 | 21355 |
| Wikipedia Categories | 3488 | 857 | 2496 | 1775 |
| Co-Author Relationship | 5727 | 4771 | 5096 | 3287 |

Table 1: Data Statistics: number of item for each field

the scholar (affiliation, full list of research interests, and similar scholars), remove the scholar from the list, and enable/disable the effect of this scholar in the final recommendation. Together with the *favorite keywords* list explained above, the list of favorite scholars form a user *profile of interests*, which the user gradually assembles while exploring possible areas of interests and scholars. In turn, the profile of interests is used to generate further recommendations as explained below.

4.1.4 Recommendations. This section (Figure 2:D) consist of two subsections. *Recommended keywords* (Figure 2:D1) shows the list of three most relevant additional keywords, which are suggested given already selected (and enabled) user's favorite keywords and scholars. User can see more information about the keyword (similar to favorite keyword section) and also add these recommended keyword to their favorite list using two circular button on the right side of each keyword. *Recommended scholars* (Figure 2:D2) shows a list of recommended scholars, which are most relevant to the active (enabled) favorite topic keywords and most similar to the active favorite scholars. For each recommended scholar the list shows basic personal and academical information. User can also see the similarity between the recommended scholars and their profile of interests represented by favorite keywords and scholars.

4.2 Recommendation method

We generate the recommendations using "Cypher Query Language" in Neo4j. In the following we explain how we generate recommendations for keywords and scholars.

4.2.1 Keyword Recommendation. In order to recommend similar keywords, we use user's favorite keywords and scholars. Each keyword is connected to other keywords in two ways: 1- Via the similar research interest between scholars and 2- Via similar relevant keywords and categories. We consider both of these relations to find similar keywords. In the final list, we sorted the keywords based on the number of occurrences then we chose top three keywords to be presented to the user.

4.2.2 Scholar Recommendation. Similar to keyword recommendation, we use both favorite keywords and scholars. There are three criteria for scholar recommendation: weighted scholar's research

interests, co-authorship relationship between scholars, and connection between scholar interest through relevant keywords and categories. After generating the list of candidate scholars, we sort it based on the similarity score (calculated based on weighted similarity score for each of three criteria) and present the top ten results to the user.

5 DISCUSSION AND FUTURE WORK

We presented a method to build a knowledge graph by integrating data from Google Scholar and Wikipedia to help students with limited knowledge about the subject find a research advisor or thesis committee member. Although Google scholar covers a variety of publications and patents, additional sources of information (scholar's active research projects, funding information, etc.) could make the knowledge graph more connected and provide the users with additional critical information when it comes to finding an advisor. We also plan to refine our keyword extraction technique. More sophisticated methods of extraction using natural language processing and machine learning could potentially improve the semantic relation between concepts and provide users with a more realistic set of research interests for scholars. We also designing a series of controlled user study and field studies to evaluate the usability and value of the exploratory search interface. We hope, that these user studies will provide valuable insights for improving the knowledge graph and the interface.

REFERENCES

- [1] Joost CF De Winter, Amir A Zadpoor, and Dimitra Dodou. 2014. The expansion of Google Scholar versus Web of Science: a longitudinal study. *Scientometrics* 98, 2 (2014), 1547–1565.
- [2] Matthew E. Falagas, Eleni I. Pitsouni, George A. Malietzis, and Georgios Pappas. 2008. Comparison of PubMed, Scopus, Web of Science, and Google Scholar: strengths and weaknesses. *The FASEB Journal* 22, 2 (2008), 338–342. <https://doi.org/10.1096/fj.07-9492LSF> arXiv:<https://doi.org/10.1096/fj.07-9492LSF> PMID: 17884971.
- [3] Philipp Mayr and Anne-Kathrin Walter. 2007. An exploratory study of Google Scholar. *Online Information Review* 31, 6 (2007), 814–830.
- [4] Olena Medelyan, David Milne, Catherine Legg, and Ian H Witten. 2009. Mining meaning from Wikipedia. *International Journal of Human-Computer Studies* 67, 9 (2009), 716–754.
- [5] John Mingers and Martin Meyer. 2017. Normalizing Google Scholar data for use in research evaluation. *Scientometrics* 112, 2 (2017), 1111–1121.
- [6] Colin Murray, Weimao Ke, and Katy Börner. 2006. Mapping Scientific Disciplines and Author Expertise Based on Personal Bibliography Files. In *Tenth International Conference on Information Visualisation (IV'06)*. 258–263. <https://doi.org/10.1109/IV.2006.73>
- [7] Enrique Orduña-Malea, Juan Manuel Ayllón, Alberto Martín-Martín, and Emilio Delgado López-Cózar. 2014. About the size of Google Scholar: playing the numbers. *arXiv preprint arXiv:1407.6239* (2014).
- [8] Behnam Rahdari and Peter Brusilovsky. 2019. User-controlled hybrid recommendation for academic papers. In *Proceedings of the 24th International Conference on Intelligent User Interfaces: Companion*. ACM, 99–100.
- [9] Richard M. Shiffrin and Katy Börner. 2004. Mapping knowledge domains. *Proceedings of the National Academy of Sciences* 101, suppl 1 (2004), 5183–5185. <https://doi.org/10.1073/pnas.0307852100>
- [10] Michael Strube and Simone Paolo Ponzetto. 2006. WikiRelate! Computing semantic relatedness using Wikipedia. In *AAAI*, Vol. 6. 1419–1424.
- [11] Max Völkel, Markus Kröttsch, Denny Vrandečić, Heiko Haller, and Rudi Studer. 2006. Semantic wikipedia. In *Proceedings of the 15th international conference on World Wide Web*. ACM, 585–594.
- [12] Wikipedia contributors. 2019. Neo4j — Wikipedia, The Free Encyclopedia. <https://en.wikipedia.org/w/index.php?title=Neo4j&oldid=892454531> [Online; accessed 2-May-2019].
- [13] Fei Wu and Daniel S. Weld. 2010. Open Information Extraction Using Wikipedia. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics (ACL '10)*. Association for Computational Linguistics, Stroudsburg, PA, USA, 118–127. <http://dl.acm.org/citation.cfm?id=1858681.1858694>